

Scalable Interpretable Multi-Response Regression via SEED *

M. Taha Bahadori, Zemin Zheng, Yan Liu and Jinchi Lv

August 15, 2016

Abstract

Sparse reduced-rank regression is an important tool to uncover meaningful dependence structure between large numbers of predictors and responses in many big data applications such as genome-wide association studies and social media analysis. Despite the recent theoretical and algorithmic advances, scalable estimation of sparse reduced-rank regression remains largely unexplored. In this paper, we suggest a scalable procedure called sequential estimation with eigen-decomposition (SEED) which needs only a single top- r singular value decomposition to find the optimal low-rank and sparse matrix by solving a sparse generalized eigenvalue problem. Our suggested method is not only scalable but also performs simultaneous dimensionality reduction and variable selection. Under some mild regularity conditions, we show that SEED enjoys nice sampling properties including consistency in estimation, rank selection, prediction, and model selection. Numerical studies on synthetic and real data sets show that SEED outperforms the state-of-the-art approaches for large-scale matrix estimation problem.

Running title: SEED

Key words: Reduced-rank regression; Scalability; Sparsity; High dimensionality; Greedy algorithm; Sparse eigenvector estimation

*Mohammad Taha Bahadori (E-mail: bahadori@gatech.edu) is Postdoctoral Fellow, School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA 30332. Zemin Zheng is Associate Professor (E-mail: zhengzm@ustc.edu.cn), Department of Statistics and Finance, School of Management, University of Science and Technology of China, China. Yan Liu is Associate Professor (E-mail: yanliu.cs@usc.edu), Computer Science Department, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089. Jinchi Lv is McAlister Associate Professor in Business Administration (E-mail: jinchilv@marshall.usc.edu), Data Sciences and Operations Department, Marshall School of Business, University of Southern California, Los Angeles, CA 90089. This work was supported by NSF CAREER Awards DMS-0955316 and IIS-1254206, Okawa Foundation Research Award, and a grant from the Simons Foundation. The authors would like to thank Emre Demirkaya and Sanjay Purushotham for their help with discussions and proofreading.

1 Introduction

Identifying complex dependence structures among predictors and responses is an important problem in statistics and machine learning, since these structures reveal hidden domain knowledge about the data. For example, in bioinformatics, identifying gene regulatory networks is crucial for understanding gene regulatory paths and gene functions, which helps disease prediction and diagnosis. Similarly, in social media analysis, inferring the influence networks from user activities (that is, *Diffusion Network Inference* problem (Leskovec et al., 2010; Zhou et al., 2013; Embar et al., 2014)) is an important problem, and it has found applications in social media marketing (Gomez-Rodriguez et al., 2012) and crisis management (Starbird and Palen, 2012). In these big data applications, inferring the dependence structures is challenging since the responses and predictors may be related through a few latent pathways and/or associated through only a subset of responses and predictors. Moreover, the curse of dimensionality and massive amounts of data, that is, scalability issues make the dependence structure discovery problem even harder to solve. Recently, regularization methods such as lasso (Tibshirani, 1996) and group lasso (Yuan and Lin, 2006), and reduced-rank regression approaches (Velu and Reinsel, 2013; Izenman, 1975) have been proposed to recover sparse response-predictor associations and latent predictors, respectively. Chen and Chan (2015) and Chen et al. (2012) have proposed sparse reduced-rank regression approaches by combining the regularization and reduced-rank regression techniques to find the complex dependence structures between responses and predictors.

Sparse reduced-rank regression works by modeling the associations between the predictor and response variables via a sparse and low-rank representation of the coefficient matrix. It not only enhances the interpretability of the estimated matrix by eliminating irrelevant features (Chen et al., 2012), but also reduces the number of free parameters of the model and thus the number of observations required for desired estimation consistency (Negahban and Wainwright, 2011; Yuan et al., 2007; Bunea et al., 2011; Candes and Plan, 2011; Chen et al., 2013). Sparse reduced-rank regression has found applications in micro-array biclustering (Chen et al., 2012), subspace clustering (Wang et al., 2013), social network community discovery (Richard et al., 2014; Zhou et al., 2013), and motion segmentation (Feng et al., 2014). In these applications, joint sparsity and low-rankness has been used to enforce a clustered dependence structure among data points. In particular, the key idea

is to estimate a similarity matrix among data points that is simultaneously sparse and low-rank and then permute the rows and columns of the matrix to yield approximately block-diagonal structures, which naturally lead to clustering of data points into several groups. Note that [Agarwal et al. \(2012\)](#), [Chandrasekaran et al. \(2010\)](#), and the references therein have considered estimating matrices with a low-rank plus sparse representation which is different from our work as we are interested in estimating a matrix that is jointly low-rank and sparse.

A natural approach to solving the sparse reduced-rank regression problem is to simultaneously penalize the parameter matrix using the L_1 and nuclear norm regularizers, as they are convex relaxations to sparsity and low-rankness of a matrix, respectively. The resulting optimization problem is convex and can be solved using the alternating direction method of multipliers (ADMM) ([Boyd et al., 2010](#)) as proposed by [Richard et al. \(2014\)](#) and [Zhou et al. \(2013\)](#). In [Bunea et al. \(2012\)](#), an alternative approach, called rank constrained group lasso (RCGL), was proposed which directly penalizes the rank and the number of nonzero rows of the parameter matrix. They showed oracle rates for the estimated matrix and also provided a practical algorithm which iteratively and jointly solves a L_1 -regularization and low-rank estimation problem. To further improve the estimation accuracy, [Chen et al. \(2012\)](#) borrowed ideas from adaptive Lasso ([Zou, 2006](#)) and proposed the iterative exclusive extraction algorithm (IEEA) which finds a locally optimal solution in the neighborhood of the initial value. They also showed model selection consistency and asymptotic normality results along with the improved empirical performance of IEEA on microarray biclustering analysis data.

All the above approaches for sparse reduced-rank regression achieve both desirable theoretical properties and strong empirical results. However, they cannot scale to large matrix estimation problems in many big data applications. The ADMM algorithm of [Richard et al. \(2014\)](#) and [Zhou et al. \(2013\)](#) uses iterative singular value thresholding ([Cai et al., 2010](#)) for solving the joint L_1 and nuclear norm regularization. Iterative singular value thresholding is known to be computationally expensive since it performs a full singular value decomposition of the parameter matrix in each iteration. On the other hand, RCGL ([Bunea et al., 2012](#)) is computationally much faster than ADMM since it only performs top- r singular value decomposition for estimating a rank- r matrix in each iteration. Despite

a lower computational cost per iteration, it is unclear how many iterations RCGL needs for convergence. IEAA (Chen et al., 2012) performs nested loops of alternating L_1 -penalized regression for each singular vector which can be expensive, especially on parallel computing devices. The iterative nature of these three approaches makes them not scalable and renders them inefficient for large matrix estimation even on high performance computing devices.

To overcome the scalability issues of the previous approaches, we propose a simple and scalable sparse reduced-rank regression procedure called sequential estimation with eigen-decomposition (SEED). SEED is designed for high-performing computing platforms. It converts the sparse and low-rank regression problem to a sparse generalized eigenvalue problem, and then solves the problem using the recent algorithms for sparse eigenvalue decomposition (Yuan and Zhang, 2013; Ma, 2013; Cai et al., 2013). As a pure learning algorithm, SEED is expected to perform only a single top- r sparse eigenvalue decomposition for estimating a rank- r matrix, which makes it truly scalable and efficient for large matrix estimation problems.

The main contributions of our paper are threefold. First, for the sparse reduced-rank regression problem, our proposed procedure SEED provides a scalable approach to uncovering the sparse predictor-response association network while simultaneously achieving dimension reduction and variable selection. Second, for the high-dimensional settings, our theoretical analysis shows that SEED can consistently estimate the singular vectors, latent factors as well as the regression coefficient matrix, identify the correct rank of the matrix, accurately predict the multivariate response vector, and recover the support of the singular vectors under mild conditions. Note that, compared to Chen et al. (2012), we do not make any assumption on the positive definiteness of the design matrix for proving our consistency results. Third, we empirically demonstrate that SEED can not only be efficiently implemented on both central processing units (CPU) and graphics processing units (GPU) for large-scale applications, but it also outperforms the state-of-the-art sparse reduced-rank regression approaches.

The rest of this paper is organized as follows. Section 2 introduces our SEED method. We discuss the implementation details of SEED in Section 3 and present its asymptotic properties in Section 4. We demonstrate the advantages of SEED on both synthetic and

real data sets in Section 5, and in Section 6 we discuss some extensions of our SEED method. All technical proofs and details are provided in the Appendix.

2 Sequential estimation with eigen-decomposition (SEED)

2.1 Data model and problem formulation

Denote by $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ n observations for the fixed design setting where $\mathbf{x}_i \in \mathbb{R}^p$ and $\mathbf{y}_i \in \mathbb{R}^q$ represent the i th predictor and the corresponding response vectors, respectively. Given a predictor vector \mathbf{x} , the corresponding response vector \mathbf{y} is drawn from the following model

$$\mathbf{y} = \mathbf{C}^{*T} \mathbf{x} + \boldsymbol{\varepsilon},$$

where the noise vector $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ is a q -dimensional zero mean multivariate Gaussian random vector with the covariance matrix $\boldsymbol{\Sigma}$, and $\mathbf{C}^* \in \mathbb{R}^{p \times q}$ is the regression coefficient matrix¹. We can rewrite the data model in the matrix form as follows

$$\mathbf{Y} = \mathbf{X}\mathbf{C}^* + \mathbf{E}, \tag{1}$$

where $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_n]^T$, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T$, and $\mathbf{E} = [\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_n]^T$ denote the matrices of stacked response, predictor and noise vectors, respectively.

Let $\mathbf{P} = n^{-1}\mathbf{X}^T\mathbf{X}$ be the Gram matrix of the predictors. We consider the case where the true regression coefficient matrix \mathbf{C}^* is jointly low-rank and sparse, the same as in Bunea et al. (2012) and Chen et al. (2012). In particular, the matrix rank r^* is assumed to be small with $r^* \ll \min(p, q)$, and \mathbf{C}^* follows a decomposition that

$$\mathbf{C}^* = \sum_{k=1}^{r^*} \mathbf{u}_k^* \mathbf{v}_k^{*T} = \sum_{k=1}^{r^*} \mathbf{C}_k^*, \tag{2}$$

where the left singular vectors $\mathbf{u}_k^* \in \mathbb{R}^p$ are \mathbf{P} -orthogonal with unit length, that is, $\mathbf{u}_k^{*T} \mathbf{P} \mathbf{u}_{k'}^* = 0$ if $k \neq k'$ and $\|\mathbf{u}_k^*\|_2 = 1$, while the right singular vectors $\mathbf{v}_k^* \in \mathbb{R}^q$ are orthogonal, that is, $\mathbf{v}_k^{*T} \mathbf{v}_{k'}^* = 0$ for $k \neq k'$, and \mathbf{C}_k^* is the layer k unit rank matrix of \mathbf{C}^* . The singular vectors

¹Note that the Gaussianity of noise variables is not essential to either our procedure or the theoretical results and similar results would hold under the sub-Gaussian assumption (Bühlmann and Van De Geer, 2011, Chapter 14).

are sorted by the magnitudes of the singular values $\sigma_k = \frac{1}{\sqrt{qn}} \|\mathbf{X}\mathbf{C}_k^*\|_F$ in descending order, which correspond to their contributions to the prediction of \mathbf{Y} .

We consider the left singular vectors (both the population and estimated ones) in the constraint space $\mathbf{u} \perp \text{Ker}(\mathbf{P})$, where $\text{Ker}(\mathbf{P})$ denotes the null space of \mathbf{P} , to guarantee the model identifiability as otherwise \mathbf{u} would contain certain component $\tilde{\mathbf{u}}$ such that $\mathbf{X}\tilde{\mathbf{u}} = 0$, which does not contribute to the prediction of \mathbf{Y} . The \mathbf{P} -orthogonality of \mathbf{u}_k^* is not necessary in the algorithm but facilitates the theoretical analysis. In fact, the above decomposition (2) for \mathbf{C}^* coincides with the singular value decomposition of $\mathbf{X}\mathbf{C}^*$ through different scalings on the singular vectors. We defer the discussion on the existence and identifiability of decomposition (2) to Supplementary Material ??.

The aforementioned modeling of the regression coefficient matrix indeed gives a latent factor model with r^* latent factors, where $\mathbf{X}\mathbf{u}_k^*$ is the k th latent factor and \mathbf{v}_k^* describes the impact of the k th factor on the response variables. As illustrated in Yuan et al. (2007), the low-rankness of \mathbf{C}^* renders dimension reduction such that all responses can be predicted by a relatively small set of common factors. Furthermore, the left singular vectors \mathbf{u}_k^* correspond to the selection of predictors, and we impose a sparsity assumption such that $\|\mathbf{u}_k^*\|_0 \ll p$ for $k = 1, \dots, r^*$. Similar sparsity assumptions were made in Bunea et al. (2012) and Chen et al. (2012) to enhance model interpretability by removing irrelevant features in high dimensions, where Chen et al. (2012) assumed that both the left and right singular vectors are sparse while Bunea et al. (2012) imposed restriction on the number of nonzero rows of the coefficient matrix. In this paper, we are interested in two cases: (i) when the right singular vectors are not required to be sparse and (ii) when it is desirable to have sparse right singular vectors, which entails the response selection. We will show that both cases are efficiently accommodated by our procedure.

Our goal is to accurately estimate the singular vectors \mathbf{u}_k^* and \mathbf{v}_k^* , and the true rank r^* such that we can recover the latent factors as well as their impacts, and at the same time, identify the underlying number of latent factors and the significant predictors. As a singular vector can have two opposite directions, we always assume that the estimated left singular vector takes the correct one, that is, the angles between estimated and population left singular vectors are no more than a right angle. Once the estimated rank and singular vectors are obtained, the estimate $\hat{\mathbf{C}}$ of the true matrix \mathbf{C}^* follows immediately from (2).

For theoretical analysis, we consider the estimation of the true rank within certain upper bound $r > r^*$, that is, $\text{rank}(\hat{\mathbf{C}}) \leq r$. In practice, the upper bound r can be controlled by our algorithm. Unlike most existing sparse and low-rank estimation methods which adopt the regularization framework of a loss function plus certain penalties and are generally not scalable, the proposed procedure SEED is indeed a pure learning algorithm that predicts \mathbf{Y} using $\mathbf{X}\hat{\mathbf{C}}$ with some low-rank $\hat{\mathbf{C}}$.

2.2 Description of SEED

The following proposition provides us insight into estimating the top- r^* left and right singular vectors of \mathbf{C}^* .

Proposition 1. *Consider the noiseless case where $\mathbf{Y}^* = \mathbf{X}\mathbf{C}^*$ and $\mathbf{C}^* = \sum_{k=1}^{r^*} \mathbf{u}_k^* \mathbf{v}_k^{*T}$ as defined in (2). Then $\mathbf{u}_1^*, \dots, \mathbf{u}_{r^*}^*$ are the r^* non-degenerate left singular vectors of \mathbf{C}^* if and only if they are the eigenvectors of the following generalized eigenvalue problem*

$$\mathbf{X}^T \mathbf{Y}^* \mathbf{Y}^{*T} \mathbf{X} \mathbf{u} = \lambda \mathbf{X}^T \mathbf{X} \mathbf{u} \quad (3)$$

with respect to the nonzero eigenvalues $\lambda_1, \dots, \lambda_{r^*}$, where $\lambda_k = nq\sigma_k^2$ is the k th largest eigenvalue of $\mathbf{Y}^* \mathbf{Y}^{*T}$ with the singular values σ_k defined in Section 2.1. Furthermore, given the left singular vector \mathbf{u}_k^* , the corresponding right singular vector \mathbf{v}_k^* can be written as

$$\mathbf{v}_k^* = \frac{1}{\mathbf{u}_k^{*T} \mathbf{X}^T \mathbf{X} \mathbf{u}_k^*} \mathbf{Y}^{*T} \mathbf{X} \mathbf{u}_k^*. \quad (4)$$

Proposition 1 shows that the problem of estimating the singular vectors can be transformed into the generalized eigenvalue problem in (3), thanks to the \mathbf{P} -orthogonality of the left singular vectors. In the noisy case where $\mathbf{Y} = \mathbf{X}\mathbf{C}^* + \mathbf{E}$, it motivates us to estimate the left singular vectors by solving the corresponding generalized eigenvalue problem as follows

$$\mathbf{X}^T \mathbf{Y} \mathbf{Y}^T \mathbf{X} \mathbf{u} = \lambda \mathbf{X}^T \mathbf{X} \mathbf{u}. \quad (5)$$

The estimation consistency in the noisy case will be guaranteed by the matrix perturbation theory, see Section 4 for details. On the other hand, it is not difficult to see that the eigenvectors with respect to different eigenvalues of problem (5) are \mathbf{P} -orthogonal, which further gives the orthogonality of the right singular vectors estimated by (4). It implies that

the right and left singular vectors obtained by solving the generalized eigenvalue problem (5) will automatically be orthogonal and \mathbf{P} -orthogonal, respectively.

Related results of principal component analysis in low dimensions can be found in Baldi and Hornik (1989), De La Torre and Black (2003), and Diamantaras and Kung (1996). Note that in the high-dimensional setting, the regime of interest for this paper, the Gram matrix \mathbf{P} can be rank deficient and the generalized eigenvalue problem is potentially challenging to solve. We will address the implementation challenges for high-dimensional settings in Section 3.

Motivated by Proposition 1, our proposed procedure SEED performs a two-step estimation for the regression coefficient matrix: it first solves the generalized eigenvalue problem in (5) to obtain the estimated left singular vectors $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_r$ with unit length; then it finds the estimated right singular vectors $\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_r$ according to (4), that is,

$$\hat{\mathbf{v}}_k = \frac{1}{\hat{\mathbf{u}}_k^T \mathbf{X}^T \mathbf{X} \hat{\mathbf{u}}_k} \mathbf{Y}^T \mathbf{X} \hat{\mathbf{u}}_k. \quad (6)$$

The maximum rank r depends on the magnitude of the estimated singular value $\hat{\sigma}_k = \frac{1}{\sqrt{qn}} \|\mathbf{X} \hat{\mathbf{C}}_k\|_F$ with $\hat{\mathbf{C}}_k = \hat{\mathbf{u}}_k \hat{\mathbf{v}}_k^T$ (whether it is larger than a threshold μ). And the optimal rank will be tuned by the information criterion described in Section 4.

The details of the procedure are provided in Algorithm SEED. To achieve a sparse solution, we need to find the optimal rank- r sparse matrix via a sparse eigenvalue decomposition procedure in Line 5 of SEED. For theoretical analysis, we assume that there exists a sparse eigenvalue decomposition procedure that solves (5) and practical methods will be provided in Section 3. The practical methods need a sparsity parameter θ , such as a threshold (Ma, 2013) or a sparsity size (Yuan and Zhang, 2013). We will show in Section 5 that SEED is robust to the choices of parameters θ and μ . If the right singular vectors are also required to be sparse, we perform a simple element-wise thresholding on $\hat{\mathbf{v}}_k$ after we obtain it in Line 7.

Given Condition 1 in Section 4, the following proposition shows that with significant probability, SEED will stop after exactly r^* iterations if the termination criterion μ is properly chosen.

Proposition 2. *Suppose that Condition 1 holds and $\log(pq) = o(n)$, then with probability at least $1 - \delta$ for any $\delta \in (0, 1)$, SEED will stop after r^* iterations if the termination criterion*

Algorithm SEED:

Input: $\mathbf{Y} \in \mathbb{R}^{n \times q}$ and $\mathbf{X} \in \mathbb{R}^{n \times p}$.

Input: A termination accuracy μ and sparsity parameter θ .

```
1 Compute matrices:  $\mathbf{P} \leftarrow \frac{1}{n} \mathbf{X}^T \mathbf{X}$ ,  $\mathbf{R} \leftarrow \frac{1}{n} \mathbf{Y}^T \mathbf{X}$ ,  $\mathbf{Q} \leftarrow \frac{1}{q} \mathbf{R}^\top \mathbf{R}$ ,  
2  $\hat{\mathbf{C}} \leftarrow \mathbf{0}_{q \times p}$   
3  $k \leftarrow 1$   
4 repeat  
5    $(\hat{\mathbf{u}}_k, \hat{\sigma}_k) \leftarrow k\text{th } \theta\text{-sparse eigenvector and eigenvalue of } \mathbf{Q}\mathbf{u} = \lambda\mathbf{P}\mathbf{u}$ .  
6   if  $\hat{\sigma}_k > \mu$  then  
7      $\hat{\mathbf{v}}_k \leftarrow \frac{1}{\hat{\mathbf{u}}_k^\top \mathbf{P} \hat{\mathbf{u}}_k} \mathbf{R} \hat{\mathbf{u}}_k$ ,      #Optional thresholding of  $\hat{\mathbf{v}}_k$  for sparsity.  
8      $\hat{\mathbf{C}} \leftarrow \hat{\mathbf{C}} + \hat{\mathbf{u}}_k \hat{\mathbf{v}}_k^\top$ .  
9      $k \leftarrow k + 1$ .  
10  end  
11 until  $\hat{\sigma}_k < \mu$   
12 return  $\hat{\mathbf{C}}$ 
```

μ is chosen from the following interval:

$$C \sqrt{\frac{1}{n} \log \frac{pq}{\delta}} \leq \mu \leq \sigma_{r^*}^2 - C \sqrt{\frac{1}{n} \log \frac{pq}{\delta}},$$

where $\sigma_{r^*}^2$ is assumed to be larger than $2C \sqrt{\frac{1}{n} \log \frac{pq}{\delta}}$ with positive constant C defined in *Supplementary Material ??*.

Proposition 2 is mainly for theoretical purposes since the interval above is generally unknown to us. We will need to tune the optimal rank by certain information criterion in practice. The tail probability δ can decay to zero quickly as p and q grow with rates such as $\delta \propto (pq)^{-\alpha}$ for some positive constant $\alpha > 1$. This is due to the fact that when $\delta \propto (pq)^{-\alpha}$, we have $\sqrt{\frac{1}{n} \log \frac{pq}{\delta}} \rightarrow 0$ under the assumption that $\log(pq) = o(n)$. The lower bound of μ excludes the case that extra latent factors are involved due to noises while the upper bound guarantees the important factors will not be missed.

3 Scalable implementation of SEED

SEED requires a sparse solution of (5) which is a generalized eigenvalue problem with a rank deficient matrix \mathbf{P} . In this section, we study multiple practical aspects of SEED and propose two different ways to solve (5), one with enhanced stability and the other using a fast procedure to accelerate it.

3.1 Basic implementation

Stability. For numerical stability purposes, we can solve the following modified problem of (5) with a very small positive ρ :

$$\mathbf{X}^T \mathbf{Y} \mathbf{Y}^T \mathbf{X} \mathbf{u} = \lambda (\mathbf{X}^T \mathbf{X} + \rho \mathbf{I}) \mathbf{u}. \quad (7)$$

Note that $\mathbf{X}^T \mathbf{X} + \rho \mathbf{I}$ is invertible since the eigenvalues of $\mathbf{X}^T \mathbf{X}$ are nonnegative. Denote by $\tilde{\mathbf{X}} \in \mathbb{R}^{p \times p}$ the modified predictor matrix such that $\tilde{\mathbf{X}}^T \tilde{\mathbf{X}} = \mathbf{X}^T \mathbf{X} + \rho \mathbf{I}$, which can be obtained via the Cholesky decomposition, and $\tilde{\mathbf{Y}} = (\tilde{\mathbf{X}}^T)^{-1} \mathbf{X}^T \mathbf{Y}$ the modified response matrix. Then the above equation (7) can be rewritten as

$$\tilde{\mathbf{X}}^T \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T \tilde{\mathbf{X}} \mathbf{u} = \lambda \tilde{\mathbf{X}}^T \tilde{\mathbf{X}} \mathbf{u}, \quad (8)$$

which adopts the same form as (5).

The formulation of $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ can be regarded as a generalization of the ridge regression to the multivariate response setting. In fact, since \mathbf{C}^* is the minimizer of $\|\mathbf{Y}^* - \mathbf{X}\mathbf{C}\|_F^2$, we can enhance the stability by adding a small Frobenius norm regularization as follows:

$$\tilde{\mathbf{C}} = \underset{\mathbf{C}}{\operatorname{argmin}} \left\{ \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \rho \|\mathbf{C}\|_F^2 \right\},$$

where the Frobenius norm is defined as $\|\mathbf{C}\|_F^2 = \sum_{i,j} \mathbf{C}_{i,j}^2$ for any matrix \mathbf{C} . After completing the squares, we get

$$\tilde{\mathbf{C}} = \underset{\mathbf{C}}{\operatorname{argmin}} \left\{ \|\tilde{\mathbf{Y}} - \tilde{\mathbf{X}}\mathbf{C}\|_F^2 \right\},$$

which means that $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ are the corresponding predictor and response matrices that take into account the shrinkage effects (James and Stein, 1961; Zheng et al., 2014).

A computationally efficient technique for solving equation (8) is to solve the sparse eigenvalue problem $\tilde{\mathbf{P}}^{-1} \tilde{\mathbf{Q}} \mathbf{u} = \lambda \mathbf{u}$, where the modified Gram matrix $\tilde{\mathbf{P}} = n^{-1} \tilde{\mathbf{X}}^T \tilde{\mathbf{X}}$ and $\tilde{\mathbf{Q}}$ is

defined accordingly as in Algorithm SEED. We also use the Sherman–Morrison–Woodbury formula to compute $\tilde{\mathbf{P}}^{-1}$ as follows:

$$(\rho \mathbf{I}_p + \mathbf{X}^T \mathbf{X})^{-1} = \frac{1}{\rho} \mathbf{I}_p - \frac{1}{\rho^2} \mathbf{X}^T (\mathbf{I}_n + \frac{1}{\rho} \mathbf{X} \mathbf{X}^T)^{-1} \mathbf{X}.$$

The above equation requires inversion of an $n \times n$ matrix instead of a $p \times p$ matrix which is significantly faster in the high-dimensional setting when $p \gg n$. The corresponding right singular vectors will then be obtained by (6).

Refitting. In SEED, a further refitting can be performed during the eigenvalue decomposition in Line 5 to enhance the stability. The refitting procedure is as follows. In the k th step, we compute the residual $\mathbf{Y}_k = \mathbf{Y} - \mathbf{X} \sum_{j=1}^{k-1} \hat{\mathbf{C}}_j$ and solve the generalized eigenvalue problem (5) with \mathbf{Y} replaced by \mathbf{Y}_k to obtain the unit rank matrix $\hat{\mathbf{C}}_k = \hat{\mathbf{u}}_k \hat{\mathbf{v}}_k^T$. Then we perform the top- k singular value decomposition $\hat{\mathbf{C}} = \mathbf{U} \hat{\mathbf{S}} \mathbf{V}^T$ for $\hat{\mathbf{C}} = \sum_{j=1}^k \hat{\mathbf{C}}_j$ and refit the solution by finding $\tilde{\mathbf{S}} = \arg\min_{\mathbf{S}} \|\mathbf{Y} - \mathbf{X} \mathbf{U} \mathbf{S} \mathbf{V}^T\|_F^2$. The estimate with refitting is defined as $\tilde{\mathbf{C}} = \mathbf{U} \tilde{\mathbf{S}} \mathbf{V}^T$. In practice, we find this approach more stable and report the results based on this variation of SEED in numerical studies.

3.2 Scalability

Speedup. The bottleneck in speeding up Algorithm SEED is Line 5 where we need to solve a sparse generalized eigenvalue problem. To overcome this bottleneck, we propose a new solution to estimating the left singular vectors by rewriting equation (5) as

$$\mathbf{X}^T (\mathbf{Y} \mathbf{Y}^T - \lambda \mathbf{I}) \mathbf{X} \mathbf{u} = \mathbf{0}.$$

Similar to Proposition 1, when \mathbf{X} is of full row rank (which is easy to satisfy in the high-dimensional setting), the above equation shares the same nonzero eigenvalues with $\mathbf{Y} \mathbf{Y}^T$. Even if \mathbf{X} is row rank deficient, the nonzero solution of \mathbf{u} is ensured by the perturbation theory in Lemma ???. Thus, we propose the following two-step procedure for Line 5 of SEED:

- (1) $\lambda \leftarrow \lambda_{\max}(\mathbf{Y} \mathbf{Y}^T)$.
- (2) $\hat{\mathbf{u}} \leftarrow$ sparse eigenvector corresponding to zero eigenvalue of $(\mathbf{X}^T \mathbf{Y} \mathbf{Y}^T \mathbf{X} - \lambda \mathbf{X}^T \mathbf{X})$.

Similarly as before, we compute the residual $\mathbf{Y}_k = \mathbf{Y} - \mathbf{X} \sum_{j=1}^{k-1} \widehat{\mathbf{C}}_j$ in the k th step and replace \mathbf{Y} with \mathbf{Y}_k in the above two-step procedure to obtain the k th left singular vector $\widehat{\mathbf{u}}_k$. Overall, the first step requires calculation of the top- r eigenvalues for an $n \times n$ matrix (or $q \times q$ if $q < n$) while the second step finds the corresponding eigenvectors by solving a regular sparse eigenvalue problem. Thus, the above procedure significantly accelerates the speed of SEED as it converts a degenerate sparse generalized eigenvalue problem to two simpler regular sparse eigenvalue problems. The truncated power method (Yuan and Zhang, 2013) can be used to compute both eigenvalue problems efficiently.

Sparse eigenvector estimation. The previous two approaches for solving the generalized eigenvalue decomposition problem indicate that we can solve the problem via regular eigenvalue decomposition. This allows us to reuse the existing procedures for sparse eigenvalue decomposition such as Yuan and Zhang (2013), Ma (2013), Cai et al. (2013), and Lei and Vu (2015) to solve the problem in (5). In numerical studies, we use the iterative thresholding method (Ma, 2013) for estimating the sparse eigenvectors of a matrix.

Parallel implementation. In order to scale up the procedures, we often need to utilize the parallel computing tools. Given the fact that SEED only uses basic matrix operations, we can employ parallel implementation of the large matrix operations to accelerate SEED. In Section 5, our experiments with GPU which contain thousands of processing units show that the matrix operations in SEED can be efficiently parallelized and it significantly enhances the speed of SEED. Whenever the data can not be loaded into the memory of a single device, efficient distributed algorithms can be used, see for example, Kang et al. (2011) and the references therein.

4 Asymptotic properties of SEED

In this section, we will analyze the statistical properties of SEED. Define the maximum sparsity level of the left singular vectors as $s^* = \max_{k=1}^{r^*} \|\mathbf{u}_k^*\|_0 \ll p$, which is assumed mainly for theoretical analysis (see Condition 1 below) and will not be used directly in the algorithm. We consider the estimated left singular vectors with the number of nonzero elements less than certain sparsity level $s > s^*$, that is, $\|\widehat{\mathbf{u}}_k\|_0 \leq s$ for $k = 1, \dots, r$.

4.1 Technical conditions

Here we list a few technical conditions and discuss their relevance in detail.

Condition 1 (Restricted isometry). *There exists a positive constant ϕ_s such that the Gram matrix \mathbf{P} satisfies*

$$\phi_s \leq \min_{\mathbf{z} \in \mathbb{R}^p} \left\{ \frac{\|\mathbf{P}\mathbf{z}\|_2}{\|\mathbf{z}\|_2} : \|\mathbf{z}\|_0 \leq 2s \right\} \leq \max_{\mathbf{z} \in \mathbb{R}^p} \left\{ \frac{\|\mathbf{P}\mathbf{z}\|_2}{\|\mathbf{z}\|_2} : \|\mathbf{z}\|_0 \leq 2s \right\} \leq \phi_s^{-1}$$

for some $s > s^*$.

Condition 2 (Minimum singular value separation). *The non-zero singular values σ_k satisfy $\sigma_k^2 - \sigma_{k+1}^2 \geq d_\sigma > 0$ for some constant d_σ and $k = 1, \dots, r^* - 1$.*

Condition 3 (Bounded eigenvalues). *The eigenvalues of the population covariance matrix of the noise vector $\boldsymbol{\varepsilon}$ satisfy $0 < \gamma_l^2 \leq \lambda_j(\boldsymbol{\Sigma}) \leq \gamma_u^2 < \infty$ for $j = 1, \dots, q$, where γ_l and γ_u are positive constants with $\gamma_u \leq c_\gamma \sigma_{r^*}$ for some positive constant c_γ .*

Condition 4 (Minimum signal strength). *There exists some positive constant $\delta \in (0, 1)$ such that the following lower bounds on the magnitudes of the non-zero elements of \mathbf{u}_k^* and \mathbf{v}_k^* hold for any $1 \leq k \leq r^*$:*

$$\begin{aligned} \min_{i \in \text{supp}(\mathbf{u}_k^*)} |u_i^*| &\geq 3C_u \sqrt{\frac{1}{n} \log \frac{pq}{\delta}}, \\ \min_{i \in \text{supp}(\mathbf{v}_k^*)} \frac{1}{\sqrt{q}} |v_i^*| &\geq 3C_v \sqrt{\frac{1}{n} \log \frac{pq}{\delta}}, \end{aligned}$$

where C_u and C_v are constants defined in Theorem 1.

Condition 1 imposes bounds on the $2s$ -sparse eigenvalues of \mathbf{P} , which is weaker than the regular bounded eigenvalue assumption since the sparse eigenvalues do not grow as fast as the regular eigenvalues when the dimensionality p grows. As a typical condition in high dimensions, it restricts the correlations between small numbers of features and thus guarantees the identifiability of the true sparse support. See, for instance, [Candes and Tao \(2005\)](#) and [Zhang \(2011\)](#) for more discussion on it.

Recall that σ_k is the k th largest singular value of $\frac{1}{\sqrt{qn}} \mathbf{X}\mathbf{C}^*$. Condition 2 requires a non-zero separation among the singular values such that the true left singular vectors are distinguishable. For ease of presentation, we assume d_σ to be a constant. In fact, d_σ can

be allowed to converge to zero asymptotically, so we indicate the effect of d_σ clearly in the constants of the theoretical results.

The elements of the unobserved noise vector $\boldsymbol{\varepsilon}$ were assumed to be independent and identically distributed (i.i.d.) in [Bunea et al. \(2012\)](#). We relax it a bit in [Condition 3](#) by imposing bounded eigenvalues for the noise covariance matrix to recover the true rank in [Theorem 2](#). Our technical argument still applies when either $\gamma_l \rightarrow 0$ or $\gamma_u \rightarrow \infty$ as long as their rates of convergence can be controlled within certain magnitudes.

The two inequalities in [Condition 4](#) are imposed for the model selection consistency of the predictors and responses, respectively. The magnitude of the minimum signal strength is $O\left(\sqrt{\frac{\log(pq)}{n}}\right)$, which is relatively mild as it converges to zero in our setting. Since \mathbf{u}_k^* is assumed to have unit length, the scaling of \mathbf{C}^* is put on \mathbf{v}_k^* such that there is an extra factor $\frac{1}{\sqrt{q}}$ in the second inequality.

4.2 Main results

Denote by $P^2 = \max_{j=1}^p \mathbf{P}_{jj}$ and $\gamma^2 = \max_{j=1}^q \boldsymbol{\Sigma}_{jj}$ the maximum diagonal component of the Gram matrix and noise covariance matrix, respectively. Without loss of generality, we assume that $V = \max_{k=1}^{r^*} \frac{1}{\sqrt{q}} \|\mathbf{v}_k^*\|_2$ is finite. Moreover, it is clear that under [Conditions 1](#) and [3](#), P and γ are also finite constants. Throughout this section, we assume that the top- r eigenvectors are obtained by solving the generalized eigenvalue problem [\(5\)](#) with the maximum sparsity level s such that [Condition 1](#) is satisfied. The estimated regression coefficient matrix is given by $\hat{\mathbf{C}} = \sum_{k=1}^{\tilde{r}} \hat{\mathbf{C}}_k$, where \tilde{r} is the optimal rank tuned by information criterion [\(9\)](#) and $\hat{\mathbf{C}}_k = \hat{\mathbf{u}}_k \hat{\mathbf{v}}_k^T$. The following theorem bounds the estimation errors of SEED with the estimated left singular vectors $\hat{\mathbf{u}}_k$ taking the correct signs as discussed before.

Theorem 1 (Estimation and prediction consistency). *Suppose that [Conditions 1](#) and [2](#) hold, γ is finite, and $\log(pq) = o(n)$, then with probability at least $1 - \delta$ for any $\delta \in (0, 1)$*

and uniformly over $k = 1, \dots, r^*$, we have

$$\begin{aligned}
\|\hat{\mathbf{u}}_k - \mathbf{u}_k^*\|_2 &\leq C_u \sqrt{\frac{1}{n} \log \frac{pq}{\delta}} + o\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right), \\
\frac{1}{\sqrt{q}} \|\hat{\mathbf{v}}_k - \mathbf{v}_k^*\|_2 &\leq C_v \sqrt{\frac{1}{n} \log \frac{pq}{\delta}} + o\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right), \\
\frac{1}{\sqrt{n}} \|\mathbf{X}(\hat{\mathbf{u}}_k - \mathbf{u}_k^*)\|_2 &\leq \phi^{-1/2} C_u \sqrt{\frac{1}{n} \log \frac{pq}{\delta}} + o\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right), \\
\frac{1}{\sqrt{q}} \|\hat{\mathbf{C}}_k - \mathbf{C}_k^*\|_F &\leq (VC_u + C_v) \left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right) + o\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right), \\
\frac{1}{\sqrt{qn}} \|\mathbf{X}(\hat{\mathbf{C}}_k - \mathbf{C}_k^*)\|_F &\leq \phi_s^{-1/2} (VC_u + C_v) \left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right) + o\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right),
\end{aligned}$$

where the constants $C_u = \frac{4\gamma P \sigma_1}{d_\sigma \phi_s^{5/2}}$ and $C_v = 2\sqrt{2} \phi_s^{-3/2} (2V \phi_s^{-1/2} + \sigma_1) C_u + 2\sqrt{2} \phi_s^{-1} \gamma P$.

Theorem 1 shows that the uniform estimation error bounds for both top- r^* singular vectors \mathbf{u}_k^* and $\frac{1}{\sqrt{q}} \mathbf{v}_k^*$, the top- r^* latent factors $\frac{1}{\sqrt{n}} \mathbf{X} \mathbf{u}_k^*$ and unit rank matrices $\frac{1}{\sqrt{q}} \mathbf{C}_k^*$, and the uniform prediction error bounds of the top- r^* latent factors are all in the same order of $O\left(\sqrt{\frac{1}{n} \log \frac{pq}{\delta}}\right)$. Similar to Proposition 2, by setting $\delta = (pq)^{-\alpha}$ with $\alpha > 1$, the tail probability will decay to zero quickly as the dimensionality p and q grow. Furthermore, the estimation and prediction accuracy would then be within the rate of $O\left(\sqrt{\frac{\log(pq)}{n}}\right)$, where the factor $\log(pq)$ reflects the curse of dimensionality as there are pq parameters in total from the regression coefficient matrix \mathbf{C}^* .

If the true rank r^* can be correctly identified, it is not difficult to see that the estimation accuracy for $\frac{1}{\sqrt{q}} \mathbf{C}^*$ will be within the rate of $O\left(\sqrt{\frac{r^* \log(pq)}{n}}\right)$ (see Corollary 1 below), which coincides with the minimax error bound for estimating the regression coefficient vector in the univariate response setting (Raskutti et al., 2011) with the dimensionality p and sparsity size s replaced by the overall dimensionality pq and true rank r^* , respectively. In view of this, the true rank r^* , instead of the maximum sparsity level s^* , plays the same role in multivariate regression as the sparsity size in the univariate response setup. With the true rank r^* , the prediction accuracy for the multivariate response vector \mathbf{y} will also follow from the prediction accuracy of the top- r^* latent factors.

Based on the discussion before, a desirable statistical property of any low-rank estimation procedure is to accurately recover the true rank of the parameter matrix. Similar to

Lasso, the nuclear norm regularization needs to be enhanced by techniques such as adaptive regularization to accurately recover the rank of the matrix (Chen et al., 2013; Bach, 2008). In contrast, in SEED we can directly control the rank of the solution by limiting the number of steps. In particular, we propose a GIC-type (Fan and Tang, 2013) information criterion that guarantees rank recovery by SEED when the optimal rank is tuned according to it.

Theorem 2 (Consistency of rank recovery). *Suppose Conditions 1–3 hold, $\log(pq) = o(n)$, $r = o\left(\left[\frac{n}{\log(pq)}\right]^{1/4}\right)$, and $r^* = o\left(\frac{1}{\sqrt{\log \log n}} \cdot \left[\frac{n}{\log(pq)}\right]^{1/4}\right)$. Define the following information criterion:*

$$\mathcal{C}_n = \sqrt{n} \log \mathcal{L}_n(\mathbf{Y}, \mathbf{X}, \hat{\mathbf{C}}) + \text{rank}(\hat{\mathbf{C}}) \sqrt{\log(pq)} \log \log n, \quad (9)$$

where $\mathcal{L}_n(\mathbf{Y}, \mathbf{X}, \mathbf{C}) = \frac{1}{qn} \|\mathbf{Y} - \mathbf{XC}\|_F^2$. Under the above information criterion, with probability at least $1 - (pq)^{-\alpha}$ for some positive constant $\alpha > 1$ and sufficiently large n , SEED will select the true rank, that is, $\text{rank}(\hat{\mathbf{C}}) = \text{rank}(\mathbf{C}^*)$.

In the high-dimensional setting where the number of predictors can increase exponentially with the sample size, it is demonstrated in Fan and Tang (2013) that we need some power of the logarithmic factor of dimensionality ($\sqrt{\log(pq)}$ for our setting) in the model complexity penalty of the information criterion to consistently identify the true model, and the slow diverging rate $\log \log n$ is set to prevent underfitting. The proof of Theorem 2 indeed shows that information criterion (9) will keep decreasing until the estimated rank reaches the true rank r^* , where in each step the amount of decrease in the objective function $\mathcal{L}_n(\mathbf{Y}, \mathbf{X}, \hat{\mathbf{C}})$ equals to the squared singular value obtained by solving the generalized eigenvalue problem (5). After reaching the true rank, the estimated singular value becomes small such that the model complexity penalty will overweight the decrease and then information criterion (9) would start increasing. Therefore, in the sequence of solutions generated by SEED, the estimate $\hat{\mathbf{C}}$ with rank r^* will be the minimizer of (9) such that the true rank can be correctly identified.

As discussed before, correct identification of the true rank will yield the estimation accuracy of \mathbf{C}^* as well as the prediction accuracy of \mathbf{XC}^* . Therefore, combined with Theorem 2, it is immediate that the results in Theorem 1 give the following corollary.

Corollary 1 (Overall estimation and prediction consistency). *Given Conditions 1–3, $\log(pq) = o(n)$, $r = o\left(\left[\frac{n}{\log(pq)}\right]^{1/4}\right)$, and $r^* = o\left(\frac{1}{\sqrt{\log \log n}} \cdot \left[\frac{n}{\log(pq)}\right]^{1/4}\right)$, if the optimal rank is tuned*

by information criterion (9), then with probability at least $1 - (pq)^{-\alpha}$ for any constant $\alpha > 0$ and sufficiently large n , we have

$$\begin{aligned} \frac{1}{\sqrt{q}} \|\hat{\mathbf{C}} - \mathbf{C}^*\|_F &\leq (1 + \alpha)(VC_u + C_v) \left(\sqrt{\frac{r^* \log(pq)}{n}} \right) + o \left(\sqrt{\frac{r^* \log(pq)}{n}} \right), \\ \frac{1}{\sqrt{qn}} \|\mathbf{X}(\hat{\mathbf{C}} - \mathbf{C}^*)\|_F &\leq \phi_s^{-1/2} (1 + \alpha)(VC_u + C_v) \left(\sqrt{\frac{r^* \log(pq)}{n}} \right) + o \left(\sqrt{\frac{r^* \log(pq)}{n}} \right). \end{aligned}$$

Besides estimation consistency and rank recovery, SEED is also able to find the true support of the singular vectors. To achieve this goal, after selecting the optimal rank, we need to further refine the model selection procedure by performing a hard-thresholding. See, for example, [Fan and Lv \(2013\)](#) for more general implications of thresholding in high-dimensional sparse modeling. Specifically, denote by $T_\theta(\mathbf{z})$ the estimator after the hard-thresholding operation on every element of $\mathbf{z} = (z_1, \dots, z_p) \in \mathbb{R}^p$ such that

$$T_\theta(z_i) = \begin{cases} 0 & \text{if } |z_i| < \theta \\ z_i & \text{otherwise} \end{cases}, \quad i = 1, \dots, p.$$

Based on the results of Theorems 1 and 2 and the signal strength assumption in Condition 4, we have the following properties for the estimator with a further thresholding.

Theorem 3 (Support recovery of the singular vectors). *Given Conditions 1–4, $\log(pq) = o(n)$, $r = o\left(\left[\frac{n}{\log(pq)}\right]^{1/4}\right)$, and $r^* = o\left(\frac{1}{\sqrt{\log \log n}} \cdot \left[\frac{n}{\log(pq)}\right]^{1/4}\right)$, for every pair of singular vectors, $(\hat{\mathbf{u}}_k, \hat{\mathbf{v}}_k)$, $k = 1, \dots, r^*$, the following results hold:*

- a) *If the threshold $\theta \in (\frac{5}{4}T_u, \frac{7}{4}T_u)$ with $T_u = C_u \sqrt{\frac{1}{n} \log \frac{pq}{\delta}}$, then with probability at least $1 - \delta$, we have $\text{supp}(T_\theta(\hat{\mathbf{u}}_k)) = \text{supp}(\mathbf{u}_k^*)$;*
- b) *If the threshold $\theta \in (\frac{5}{4}T_v, \frac{7}{4}T_v)$ with $T_v = C_v \sqrt{\frac{1}{n} \log \frac{pq}{\delta}}$, then with probability at least $1 - \delta$, we have $\text{supp}(T_\theta(\hat{\mathbf{v}}_k)) = \text{supp}(\mathbf{v}_k^*)$.*

Theorem 3 shows that both supports of the left and right singular vectors can be accurately recovered with properly chosen tuning parameter θ . Together with the correctly identified true rank r^* , the above results indeed yield consistent selection of both predictors and responses. In practice, this threshold θ can be tuned by criteria such as cross-validation.

Our simulation studies in Section 5 show that SEED is robust to the choice of tuning parameter θ .

Besides the statistical properties established before, the proposed procedure SEED enjoys great flexibility in the sense that it does not rely on exact eigenvalue decomposition and the perturbation errors in the generalized eigenvalue problem (5) will be linearly incorporated into the estimated singular vectors $\hat{\mathbf{u}}_k$ and $\hat{\mathbf{v}}_k$. Furthermore, our analysis does not rely on the positive definiteness of the Gram matrix (Chen et al., 2012) in high dimensions.

5 Numerical studies

In this section, we conduct experiments on three data sets, including two simulation data sets (one for a medium-scale experiment and one for a large-scale experiment) and one application data set in social media analysis, to examine the empirical performance of SEED.

5.1 Simulation studies

5.1.1 Simulation example 1

We generate a medium-scale synthetic data set as follows: the predictors \mathbf{x} are drawn from a multivariate Gaussian distribution as $\mathbf{x} \sim \mathcal{N}(\mathbf{0}_{p \times 1}, \mathbf{\Sigma}_X)$, where $\mathbf{\Sigma}_X$ is the $p \times p$ covariance matrix with auto-regressive structure, that is, $\Sigma_{X,i,j} = \rho^{|i-j|}$ for some $0 < \rho < 1$ which will be specified later. The responses \mathbf{y} are drawn according to conditional distribution $\mathbf{y}|\mathbf{x} \sim \mathcal{N}(\mathbf{C}^\top \mathbf{x}, \gamma \mathbf{\Sigma}_E)$ where the noise covariance matrix $\mathbf{\Sigma}_E$ is also selected to have the autoregressive structure with $\rho = 0.5$ and we set $\gamma = 0.1$. We generate the parameter matrix \mathbf{C} as follows: first we generate a block-sparse matrix $\tilde{\mathbf{C}}$ with 5% non-zero elements. Each non-zero element of $\tilde{\mathbf{C}}$ is drawn from a $\mathcal{N}(0, 1)$. To achieve a low-rank structure, we find the top- r singular value decomposition of $\tilde{\mathbf{C}}$ as $\tilde{\mathbf{C}} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$, and then set the elements of \mathbf{U} and \mathbf{V} whose magnitude is smaller than 0.01 to zero to obtain $\bar{\mathbf{U}}$ and $\bar{\mathbf{V}}$. The final parameter matrix is obtained as $\mathbf{C} = \bar{\mathbf{U}}\bar{\mathbf{S}}\bar{\mathbf{V}}^\top$ where the first r diagonal elements of the diagonal matrix $\bar{\mathbf{S}}$ are set to 100, 99, \dots , $101 - r$. Without loss of generality, we add a few vectors to the design matrix to ensure the orthogonality condition of Section 4.1. In all of the simulation experiments, we generate 100 data sets and report the mean and standard

error of performance for different methods.

We compare the performance of SEED with two state-of-art methods: (1) RCGL (Bunea et al., 2012) and (2) Penalized regression with simultaneous L_1 and nuclear-norm penalization. The optimization problem is solved by the popular alternating direction method of multipliers (Boyd et al., 2010) and we will refer to this baseline as the “LN-ADMM” algorithm. All model parameters are set based on a separate validation set with size $n_{\text{valid}} = 500$.

The quality of the estimator $\hat{\mathbf{C}}$ is evaluated via four performance metrics listed as follows. (1) *Normalized Prediction Error* defined as:

$$\text{Normalized Prediction Error} = \frac{\|\mathbf{Y}_{\text{test}} - \mathbf{X}_{\text{test}}\hat{\mathbf{C}}\|_F}{\|\mathbf{Y}_{\text{test}}\|_F}.$$

(2) *Normalized Parameter Estimation Error* defined as:

$$\text{Normalized Parameter Estimation Error} = \frac{\|\hat{\mathbf{C}} - \mathbf{C}\|_F}{\|\mathbf{C}\|_F},$$

where \mathbf{C} is the true parameter matrix.

(3) *Rank Recovery Error* defined as:

$$\text{Rank Recovery Error} = |\text{rank}(\hat{\mathbf{C}}) - \text{rank}(\mathbf{C})|.$$

Since the solution of the nuclear norm always leads to small non-zero singular values (which prevents $\hat{\mathbf{C}}$ from being low-rank), we threshold the singular values of $\hat{\mathbf{C}}$ that are more than 100 times smaller than its largest singular value to have a fair comparison.

(4) *Support Recovery AUC*, that is, the area under the receiver operating characteristic (ROC) curve of comparing support of $\hat{\mathbf{C}}$ with the ground truth, which is always between 0 and 1. It is computed by varying the decision threshold and obtaining the false positive and true positive curve. Then the area under the false positive and true positive curve is reported as AUC. The value of AUC indicates the probability that a procedure assigns a higher value to a randomly chosen non-zero element than a randomly chosen zero element (Hanley and McNeil, 1982). It is an appropriate metric for measuring support recovery accuracy because in sparse support recovery we have more zeros than non-zeros which inflates the result of the simple 0-1 accuracy measure. In contrast, AUC is more robust to imbalanced positive/negative prediction labels.

Table 1 shows the results of all algorithms on a variety of regimes by varying the dimensionality p and the rank r . We can see that SEED is superior or comparable to the




Figure 1: (a) Solution path for the singular values of the estimated matrices. The plot show the value of top five singular values of the solution $\hat{\mathbf{C}}$ as we change the stopping error μ . (b) Solution path for the top left singular vector \mathbf{u} of the estimated matrices. Only seven coefficients are non-zero. The range of the parameters are generated as follows: $\mu = \text{logspace}(-5, -1, 5)$ and $\theta = \text{logspace}(-1, \log_{10}(20), 10)$, where $\text{logspace}(a, b, n)$ indicates the minimum value 10^a , maximum value 10^b , and total number n .

baseline algorithms across all four measures. As the results show and the theory predicts, in most high-dimensional cases, nuclear norm usually overestimates the true rank of the matrix. Furthermore, we find that the iterative SVD procedure in the RCGL algorithm often results in significant underestimation of the true rank, when the true rank is large. Note that in addition to accuracy, SEED also significantly reduces the variance of the estimation.

Figure 1a shows the solution path for SEED on one example data set ($p = 400$, $r = 5$, $q = 200$, $\rho = 0.5$, and $n = 100$). The corresponding singular values are set to 30, 27, 24, 21, and 18. In the horizontal axis, we show the termination parameter μ normalized by $\|\mathbf{Y}\|_F^2/(nq)$. We can see that SEED can identify the correct rank with medium values of μ . Figure 1b shows the solution path for the top left singular vector \mathbf{u}_1 of \mathbf{C} on an example data set ($p = 200$, $q = 100$, $n = 50$, $\rho = 0.5$, and $r = 1$). Both of the solution paths indicate that SEED is robust to the particular choice of parameters and in a large range of parameters SEED is able to successfully recover the true rank of the matrix and the support of the singular vectors.

Table 1: Simulation Results

p	Algorithm	Normalized Prediction Error ($\times 10^{-2}$)	Normalized Estimation Error ($\times 10^{-2}$)	Rank Recovery Error	Support Recovery AUC
$n = 100, q = 200, \mathbf{r} = \mathbf{3}, \text{ and } \rho = 0.5$					
100	SEED	1.8958 (0.0587)	0.5762 (0.0312)	0.0000 (0.0000)	0.9816 (0.0026)
	LN-ADMM	3.1031 (0.0529)	12.4132 (0.1736)	0.6100 (0.0634)	0.8078 (0.0020)
	RCGL	4.3006 (0.0619)	16.3268 (0.1708)	0.7600 (0.0698)	0.8024 (0.0019)
400	SEED	0.8140 (0.0121)	0.3568 (0.0069)	0.0000 (0.0000)	0.9706 (0.0017)
	LN-ADMM	13.5766 (0.1356)	27.3787 (0.1328)	1.6700 (0.0711)	0.8085 (0.0050)
	RCGL	29.7221 (0.2244)	34.6811 (0.2674)	0.8400 (0.0692)	0.7174 (0.0029)
800	SEED	0.6296 (0.0084)	0.3603 (0.0063)	0.0000 (0.0000)	0.9720 (0.0012)
	LN-ADMM	5.6560 (0.0504)	15.8111 (0.0774)	0.3500 (0.0500)	0.9948 (0.0001)
	RCGL	6.5201 (0.0419)	19.4005 (0.0847)	0.4700 (0.0627)	0.9911 (0.0001)
1500	SEED	0.4891 (0.0048)	0.3578 (0.0052)	0.0000 (0.0000)	0.9718 (0.0011)
	LN-ADMM	5.4634 (0.0238)	22.6509 (0.0533)	3.0700 (0.0624)	0.9949 (0.0001)
	RCGL	8.2969 (0.0481)	29.3042 (0.0780)	1.3700 (0.2159)	0.9871 (0.0001)
2000	SEED	0.4351 (0.0049)	0.3546 (0.0060)	0.0000 (0.0000)	0.9738 (0.0009)
	LN-ADMM	4.9001 (0.0224)	23.1360 (0.0656)	2.8700 (0.0812)	0.9962 (0.0001)
	RCGL	8.7672 (0.0501)	32.0000 (0.0897)	2.5100 (0.3043)	0.9879 (0.0001)
$n = 100, q = 200, \mathbf{r} = \mathbf{30}, \text{ and } \rho = 0.5$					
100	SEED	0.5820 (0.0058)	2.2812 (0.0144)	0.0000 (0.0000)	0.8132 (0.0022)
	LN-ADMM	3.1031 (0.0529)	12.4132 (0.1736)	0.6100 (0.0634)	0.8078 (0.0020)
	RCGL	4.3006 (0.0619)	16.3268 (0.1708)	0.7600 (0.0698)	0.8024 (0.0019)
400	SEED	1.6704 (0.1641)	7.1831 (0.6961)	0.0000 (0.0000)	0.8053 (0.0044)
	LN-ADMM	13.5766 (0.1356)	27.3787 (0.1328)	1.6700 (0.0711)	0.8085 (0.0050)
	RCGL	29.7221 (0.2244)	34.6811 (0.2674)	0.8400 (0.0692)	0.7174 (0.0029)
800	SEED	0.1721 (0.0014)	0.4592 (0.0046)	0.0000 (0.0000)	0.9993 (0.0000)
	LN-ADMM	5.6560 (0.0504)	15.8111 (0.0774)	0.3500 (0.0500)	0.9948 (0.0001)
	RCGL	6.5201 (0.0419)	19.4005 (0.0847)	0.4700 (0.0627)	0.9911 (0.0001)
1500	SEED	0.1299 (0.0008)	0.4440 (0.0031)	0.0000 (0.0000)	0.9992 (0.0000)
	LN-ADMM	5.4634 (0.0238)	22.6509 (0.0533)	3.0700 (0.0624)	0.9949 (0.0001)
	RCGL	8.2969 (0.0481)	29.3042 (0.0780)	1.3700 (0.2159)	0.9871 (0.0001)
2000	SEED	0.1130 (0.0007)	0.4380 (0.0032)	0.0000 (0.0000)	0.9992 (0.0000)
	LN-ADMM	4.9001 (0.0224)	23.1360 (0.0656)	2.8700 (0.0812)	0.9962 (0.0001)
	RCGL	8.7672 (0.0501)	32.0000 (0.0897)	2.5100 (0.3043)	0.9879 (0.0001)




Figure 2: Speedup by SEED on (a) CPU and (b) GPU devices. Note that the vertical axis is in logarithmic scale.

5.1.2 Simulation example 2

In order to study scalability of SEED, we conduct the experiments on two computing environment, including: (1) an off-the-shelf personal computer (PC) and (2) a graphics processing unit (GPU), to demonstrate the runtime efficiency and the parallelization capability of SEED.

First, we run our experiments on an off-the-shelf PC with Intel i7 at 3.4GHz and 8GB of memory. The system runs MATLAB R2013b on the Windows operating system. We generate 5 data sets with $r = 1$, non-zero ratio of 10%, $q = 1000$, and $n = 1000$. Figure 2a shows the average CPU runtime of three algorithms as the dimension p increases. We can see that SEED can achieve a speed up of 10-100 times in runtime compared with baseline methods.

Next, in order to test scalability of SEED in extremely large data sets, we use a machine that is equipped with a Tesla K40 GPU which has 2880 processing cores at 745MHz and 12GB of memory. We perform our experiments with MATLAB R2013b on a Debian Linux operating system. GPUs are built to have many less-powerful processing units which makes them ideal for parallel implementation (Bekkerman et al., 2012, Chapter 5). Therefore we apply the two-step fast eigenvalue decomposition described in Section 3, which involves only simple matrix operation and can be paralleled easily. The experiment results shown in Figures 2b and 3 are obtained under the setting of $q = 10000$, $n = 5000$, $r = 1$ and




Figure 3: Scalability experiments on very large data sets on GPU using the fast approach. Accuracy results are normalized.

non-zero ratio of 10%. The results indicate that while SEED is fast on the GPU, it also achieves reasonable accuracy. Note that the results show that SEED is able to estimate a sparse and low-rank matrix with 10^8 elements in less than a minute, confirming its extreme scalability.

5.2 Real data analysis

Diffusion Network Inference, that is, the task of inferring influence networks from user activities, is one of the central tasks in social networks analysis (Gomez-Rodriguez et al., 2012) because it helps improve social marketing by finding the influential users in a network. It is a challenging problem because: (i) in many social networks the influence is expressed implicitly (Gomez-Rodriguez et al., 2012) and (ii) empirical studies show that common metrics such as number of friends or followers fail to accurately measure the social influence of the users (Cha et al., 2010).

A popular computational approach in estimating social influence among users is to count the number of users' activities over a time span (in regularly or irregularly spaced intervals) and analyze the resulting time series data (Truccolo et al., 2005). Many different models have been developed, among which the vector auto-regressive model arises as a simple and robust solution (Trusov et al., 2009; Bahadori and Liu, 2013). That is, every user is described by a time series $x_i(t)$ for $t = 1, \dots, T$. Next, the vector auto-regressive

model with L lags is fitted to the time series as follows:

$$\mathbf{x}(t) = \sum_{\ell=1}^L A_{\ell} \mathbf{x}(t - \ell),$$

where $\mathbf{x}(t) = [x_1(t), \dots, x_p(t)]$ and A_{ℓ} is the evolution matrix at ℓ th lag. The influence network can be built from the evolution matrices by establishing an edge from node j to node i if $\sum_{\ell=1}^L |A_{\ell, i, j}|$ is significantly larger than zero.

In this experiment, we gather a Twitter data set with tweets on the ‘‘Haiti earthquake’’ and apply vector auto-regressive model to identify the potential top influencers on this topic (that is, those Twitter accounts with the largest impact on the others). We divide the 17 days after the Haiti Earthquake on Jan. 12, 2010 into 1000 uniformly spaced intervals and generate a multivariate time series data set by counting the number of tweets on this topic for the top 1000 users who tweeted most about it. For accurate modeling, we remove the users that were highly correlated with each other, most of which were operated by the same users and tweeted exactly the same content. We also remove robot-like user-accounts which tweeted on very regular intervals, which in total led to a subset of 270 users. We analyze this data with a VAR(5) model which requires estimation of a $q = 270$ dimensional response vector using $p = 1350$ predictors while we have only $n = 995$ observations. The number of lags is chosen based on the intuition about the maximum retweeting delay.

Since we do not have access to the true influence network, we use the retweet network as a surrogate of the ground truth following the evaluation convention in the social networks community. The retweet network is constructed by adding an edge from user i to user j if user j has retweeted at least 4 of the tweets of user i . Clearly, the retweet network is not the actual underlying temporal dependency graph, mainly because there are possible implicit influence patterns as well. However, it is the best possible metric that we could obtain for graph estimation accuracy evaluation in our data set (Cha et al., 2010). The retweet network for the 270 selected users is sparse; it has only 0.11% of possible edges.

We apply SEED, LN-ADMM, and RCGL algorithms to uncover the influence network in our twitter data set. Figure 4 shows the accuracy of the procedures in uncovering the true influence network in terms of AUC. For every value of the rank parameter, we tune the sparsity by 5-fold cross-validation. Given the fact that exact rank constraint cannot be enforced directly in the LN-ADMM algorithm, we find the best value of the nuclear norm




Figure 4: The graph recovery accuracy of the algorithms as the rank of solution varies.

Table 2: Run time (in seconds) of the algorithms on the application data set.

SEED	LN-ADMM	RCGL
2.83	127.34	256.87

regularization parameter λ_L by 5-fold cross-validation. Then, we compute the low-rank approximations of the parameter matrix and evaluate the accuracy at each rank.

The results in Figure 4 show that SEED significantly outperforms the baseline procedures. They also indicate that, in all of the algorithms, as we increase the rank of the solution matrix, the accuracy is improved initially and then quickly saturates. SEED achieves the highest accuracy when the rank is 4. Note that this result also confirms other studies that the social network connections may be strongly influenced by a few unobserved exogenous variables (Myers et al., 2012). The results in Table 2 demonstrate the significant speedup achieved by SEED compared to the baselines.

6 Discussion

In this paper, we propose to convert the problem of sparse reduced-rank regression into a sparse generalized eigenvalue problem, which allows us to efficiently employ the recently developed sparse eigenvalue decomposition techniques. After this transformation, the left singular vectors can be estimated in two simple steps, and the estimation of both sparse and dense right singular vectors is unified in a single framework. As a pure learning algorithm, SEED deviates from traditional regularization frameworks (i.e., a loss function plus certain

penalties), leading to computational efficiency and scalability. Furthermore, we prove that SEED achieves nice estimation and prediction accuracy that coincides with the minimax error bound in the univariate regression setting (Raskutti et al., 2011).

Some interesting problems for future research include extending the current formulation of the regression coefficient matrix in (2) to the case where the singular values can be repeated such that the left singular vectors (which correspond to latent factors) are not identifiable. Then we will need to estimate the eigenspaces spanned by important singular vectors and characterize the estimation accuracy by some new criterion, such as the one in Cai et al. (2013) and Ma (2013). Another research direction is to explore the theory of random design matrices and this can be addressed by using an extended version of perturbation theory where the perturbation in \mathbf{P} is also included in the analysis.

Moreover, it is computationally straightforward to extend SEED to the generalized linear models by adapting the sequential quadratic programming framework. For this extension, we first approximate the loss function by the quadratic loss function and find the optimal unit rank matrix. Then we can add the unit rank matrix to the solution and re-approximate the loss function with another quadratic function around this new solution. By performing these three steps sequentially, we can efficiently estimate the low-rank coefficient matrix. Statistical properties of such estimator can be analyzed by extending the results in Lozano et al. (2011) for greedy sparse procedures to reduced-rank regression.

References

- Agarwal, A., Negahban, S., Wainwright, M. J., et al. (2012). Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*, 40(2):1171–1197.
- Bach, F. R. (2008). Consistency of trace norm minimization. *The Journal of Machine Learning Research*, 9:1019–1048.
- Bahadori, M. T. and Liu, Y. (2013). An Examination of Practical Granger Causality Inference. In *SIAM Data Mining*, pages 467–475.

- Baldi, P. and Hornik, K. (1989). Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58.
- Bekkerman, R., Bilenko, M., and Langford, J. (2012). *Scaling up machine learning: Parallel and distributed approaches*. Cambridge University Press.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2010). Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.
- Bühlmann, P. and Van De Geer, S. (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media.
- Bunea, F., She, Y., and Wegkamp, M. H. (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics*, 39(2):1282–1309.
- Bunea, F., She, Y., Wegkamp, M. H., et al. (2012). Joint variable and rank selection for parsimonious estimation of high-dimensional matrices. *The Annals of Statistics*, 40(5):2359–2388.
- Cai, J.-F., Candès, E. J., and Shen, Z. (2010). A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982.
- Cai, T. T., Ma, Z., and Wu, Y. (2013). Sparse pca: Optimal rates and adaptive estimation. *The Annals of Statistics*, 41(6):3074–3110.
- Candes, E. J. and Plan, Y. (2011). Tight Oracle Inequalities for Low-Rank Matrix Recovery From a Minimal Number of Noisy Random Measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359.
- Candes, E. J. and Tao, T. (2005). Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215.
- Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, P. K. (2010). Measuring User Influence in Twitter: The Million Follower Fallacy. *ICWSM*, 10:10–17.

- Chandrasekaran, V., Parrilo, P., Willsky, A. S., et al. (2010). Latent variable graphical model selection via convex optimization. In *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, pages 1610–1613. IEEE.
- Chen, K. and Chan, K.-S. (2015). A note on rank reduction in sparse multivariate regression. *Journal of Statistical Theory and Practice*, 10(1):100–120.
- Chen, K., Chan, K.-S., and Stenseth, N. C. (2012). Reduced rank stochastic regression with a sparse singular value decomposition. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(2):203–221.
- Chen, K., Dong, H., and Chan, K.-S. (2013). Reduced rank regression via adaptive nuclear norm penalization. *Biometrika*, 100(4):901–920.
- De La Torre, F. and Black, M. J. (2003). A framework for robust subspace learning. *International Journal of Computer Vision*, 54(1-3):117–142.
- Diamantaras, K. I. and Kung, S. Y. (1996). *Principal component neural networks: theory and applications*. John Wiley & Sons, Inc.
- Embar, V. R., Pasumarthi, R. K., and Bhattacharya, I. (2014). A bayesian framework for estimating properties of network diffusions. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1216–1225. ACM.
- Fan, Y. and Lv, J. (2013). Asymptotic equivalence of regularization methods in thresholded parameter space. *Journal of the American Statistical Association*, 108:1044–1061.
- Fan, Y. and Tang, C. Y. (2013). Tuning parameter selection in high dimensional penalized likelihood. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3):531–552.
- Feng, J., Lin, Z., Xu, H., and Yan, S. (2014). Robust Subspace Segmentation with Block-Diagonal Prior. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 3818–3825. IEEE.
- Gomez-Rodriguez, M., Leskovec, J., and Krause, A. (2012). Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):21.

- Hanley, J. A. and McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36.
- Izenman, A. J. (1975). Reduced-rank regression for the multivariate linear model. *Journal of multivariate analysis*, 5(2):248–264.
- James, W. and Stein, C. (1961). Estimation with quadratic loss. In *Proc. 4th Berkeley Symp. Mathematical Statistics and Probability*, volume 1, pages 361–379. Berkeley: University of California Press.
- Kang, U., Meeder, B., and Faloutsos, C. (2011). Spectral Analysis for Billion-Scale Graphs: Discoveries and Implementation. In *Advances in Knowledge Discovery and Data Mining*, volume 6635, pages 13–25.
- Lei, J. and Vu, V. Q. (2015). Sparsistency and agnostic inference in sparse pca. *The Annals of Statistics*, 43(1):299–322.
- Leskovec, J., Chakrabarti, D., Kleinberg, J., Faloutsos, C., and Ghahramani, Z. (2010). Kronecker Graphs: An Approach to Modeling Networks. *Journal of Machine Learning Research*, 11:985–1042.
- Lozano, A. C., Swirszcz, G., and Abe, N. (2011). Group orthogonal matching pursuit for logistic regression. In *International Conference on Artificial Intelligence and Statistics*, pages 452–460.
- Ma, Z. (2013). Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*, 41(2):772–801.
- Myers, S. A., Zhu, C., and Leskovec, J. (2012). Information diffusion and external influence in networks. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 33–41. ACM.
- Negahban, S. and Wainwright, M. J. (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics*, 39(2):1069–1097.
- Raskutti, G., Wainwright, M., and Yu, B. (2011). Minimax rates of estimation for high-dimensional linear regression over ℓ_q balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994.

- Richard, E., Gaïffas, S., and Vayatis, N. (2014). Link Prediction in Graphs with Autoregressive Features. *Journal of Machine Learning Research*, 15:565–593.
- Starbird, K. and Palen, L. (2012). (how) will the revolution be retweeted?: information diffusion and the 2011 egyptian uprising. In *Proceedings of the acm 2012 conference on computer supported cooperative work*, pages 7–16. ACM.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 267–288.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of neurophysiology*, 93(2):1074–89.
- Trusov, M., Bucklin, R. E., and Pauwels, K. (2009). Effects of word-of-mouth versus traditional marketing: findings from an internet social networking site. *Journal of marketing*, 73(5):90–102.
- Velu, R. and Reinsel, G. C. (2013). *Multivariate reduced-rank regression: theory and applications*, volume 136. Springer Science & Business Media.
- Wang, Y.-X., Xu, H., and Leng, C. (2013). Provable subspace clustering: When LRR meets SSC. In *Advances in Neural Information Processing Systems*, pages 64–72.
- Yuan, M., Ekici, A., Lu, Z., and Monteiro, R. (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(3):329–346.
- Yuan, M. and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67.
- Yuan, X.-T. and Zhang, T. (2013). Truncated Power Method for Sparse Eigenvalue Problems. *Journal of Machine Learning Research*, 14(1):899–925.
- Zhang, T. (2011). Adaptive Forward-Backward Greedy Algorithm for Learning Sparse Representations. *IEEE Transactions on Information Theory*, 57(7):4689 – 4708.

- Zheng, Z., Fan, Y., and Lv, J. (2014). High dimensional thresholded regression and shrinkage effect. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(3):627–649.
- Zhou, K., Zha, H., and Song, L. (2013). Learning Social Infectivity in Sparse Low-rank Networks Using Multi-dimensional Hawkes Processes. In *Proceedings of International Conference on Artificial Intelligence and Statistics*, volume 31, pages 641–649.
- Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476):1418–1429.